Network Working Group
Request for Comments: 932

David D. Clark
MIT, LCS
January 1985

Postel (Chone)

A SUBNETWORK ADDRESSING SCHEME

STATUS OF THIS MEMO

This RFC suggests a proposed protocol for the ARPA-Internet community, and requests discussion and suggestions for improvements. Distribution of this memo is unlimited.

INTRODUCTION

Several recent RFCs have discussed the need for a "subnet" structure within the internet addressing scheme, and have proposed strategies for "subnetwork" addressing and routing. In particular, Jeff Mogul in his RFC-917, "Internet Subnets", describes an addressing scheme in which a variable number of the leading bits of the host portion of the address are used to identify the subnet. The drawback to this scheme is that it is necessary to modify the host implementation in order to implement it. While the modification is a simple one, it is necessary to retrofit it into all implementations, including those which are already in the field. (See RFC-917 by Mogul for various alternative approaches to this problem, such as using Address Resolution Protocol.)

This RFC proposes an alternative addressing scheme for subnets which, in most cases, requires no modification to host software whatsoever. The drawbacks of this scheme are that the total number of subnets in any one network are limited, and that modification is required to all gateways.

THE PROPOSAL

In this scheme, the individual subnets of a network are numbered using Class C addresses. Since it is necessary with this scheme that a Class C address used to number a subnet be distinguishable from a Class C address used to number an isolated network, we will reserve for subnetworks the upper half of the Class C address space, in other words all those Class C addresses for which the high order bit is on. When a network is to be organized as a series of subnetworks, a block of these reserved Class C addresses will be assigned to that network, specifically a block of 256 addresses having the two first bytes identical. Thus, the various subnetworks of a network are distinguished by the third byte of the Internet address. (This addressing scheme implies the limitation that there can only be 256 subnetworks in a net. If more networks are required, two blocks will have to be allocated, and the total viewed as two separate networks.)

The gateways and hosts attached to this subnetted network use these addresses as ordinary Class C addresses. Thus, no modification to any host software is required for hosts attached to a subnetwork.

For gateways not directly attached to the subnetted network, it is an unacceptable burden to separately store the routing information to each of the subnets. The goal of any subnet addressing scheme is to provide a strategy by which distant gateways can store routing information for the network as a whole. In this scheme, since the first two bytes of the address is the same for every subnet in the network, those first two bytes can be stored and manipulated as if they are a single Class B address by a distant gateway. These addresses, which can be used either as a Class B or Class C address as appropriate, have been informally called Class "B 1/2" addresses.

In more detail, a gateway would treat Class C addresses as follows under the scheme. First, test to see whether the high order bit of the address is on. If not, the address is an ordinary Class C address and should be treated as such.

If the bit is on, this Class C address identifies a subnet of a network. Test to see if this gateway is attached to that network. If so, treat the address as an ordinary Class C address.

If the gateway is not attached to the network containing that subnetwork, discard the third byte of the Class C address and treat the resulting two bytes as a Class B address. Note that there can be no conflict between this two-byte pattern and an ordinary Class B address, because the first bits of this address are not those of a valid Class B address, but rather those of a Class C address.

OPTIMIZATIONS

If a network grows to more than 256 subnetworks, it will be necessary to design two distinct blocks of special Class C addresses, and to view this aggregate as two separate networks. However, the gateways of these two networks can, by proper design, run a joint routing algorithm which maintains optimal routes between the two halves, even if they are connected together by a number of gateways.

Indeed, in general it is possible for gateways that are not directly attached to a subnetworked network to be specially programmed to remember the individual Class C addresses, if doing so provides greatly improved network efficiency in some particular case.

It was stated earlier that no modification to the host software is necessary to implement this scheme. There is one case in which a

minor modification may prove helpful. Consider the case of a distant host, not immediately attached to this subnetworked network. That host, even though at a distance, will nonetheless maintain separate routing entries for each of the distinct subnetwork addresses about which it has any knowledge. For most hosts, storing this information for each subnet represents no problem, because most implementations do not try to remember routing information about every network address in the Internet, but only those addresses that are of current interest. If, however, for some reason the host has a table which attempts to remember routing information about every Internet address it has ever seen, than that host should be programmed to understand the gateway's algorithm for collapsing the addresses of distant subnets from three bytes to two. However, it is not a recommended implementation strategy for the host to maintain this degree of routing information, so under normal circumstances, the host need not be concerned with the C to B conversion.

DRAWBACK

The major drawback of this scheme is that any implementation storing large tables of addresses must be changed to know the "B 1/2" conversion rule. Most importantly, all gateways must be programmed to know this rule. Thus, adoption of this scheme will require a scheduled mandatory change by every gateway implementation. The difficulty of organizing this is unknown.

OTHER VARIATIONS

It is possible to imagine other variations on the patterns of collapsing addresses. For example, 256 Class B addresses could be gathered together and collapsed into one Class A address. However, since the first three bits of the resulting Class A address would be constrained, this would permit only 32 such subnetted networks to exist. A more interesting alternative would be to permit the collapse of Class C addresses into a single Class A address. It is not entirely obvious the best way of organizing the sub-fields of this address, but this combination would permit a few very large nets of subnets to be assembled within the Internet.

The most interesting variation of "B 1/2" addresses is to increase the number of bits used to identify the subnet by taking bits from the resulting Class B address. For example, if 10 bits were used to identify the subnet (providing 1024 subnets per network), then the gateway, when forming the equivalent address, would not only drop the third byte but also mask the last two bits of the B address. Since the first three bits of the address are constrained, this would leave 13 bits for the network number, or 8192 possible subnetworked

networks. This number is not as large as would be desirable, so it is clear that selecting the size of the subnet field is an important compromise.

Danny Cohen has suggested that this scheme should be fully generalized so that the boundaries between the network, subnetwork, and host field be arbitrarily movable. The problem in such a generalization is to determine how the gateway is to maintain the table or algorithm which permits the collapsing of the address to occur. This RFC proposes that, in the short run, only one single form of "B 1/2" addresses be implemented as an Internet subnet standard.